

Cooperative coevolution of real predator robots and virtual robots in the pursuit domain[☆]

Lijun Sun^{a,b}, Chao Lyu^{c,a}, Yuhui Shi^{a,*}

^a*Shenzhen Key Laboratory of Computational Intelligence, Department of Computer Science and Engineering, Southern University of Science and Technology, China*

^b*Centre for Artificial Intelligence, CIBCI Lab, Faculty of Engineering and Information Technology, University of Technology Sydney, Australia*

^c*Harbin Institute of Technology, China*

Abstract

The pursuit domain, or predator-prey problem is a standard testbed for the study of coordination techniques. In spite that its problem setup is apparently simple, it is challenging for the research of the emerged swarm intelligence. This paper presents a particle swarm optimization (PSO) based cooperative coevolutionary algorithm for the (predator) robots, called CCPSO-R, where real and virtual robots coexist in an evolutionary algorithm (EA). Virtual robots sample and explore the vicinity of the corresponding real robots and act as their action spaces, while the real robots consist of the real predators who actually pursue the prey robot without fixed behavior rules under the immediate guidance of the fitness function, which is designed in a modular manner with very limited domain knowledge. In addition, kinematic limits and collision avoidance considerations are integrated into the update rules of robots. Experiments are conducted on a scalable swarm of predator robots with 4 types of preys, the results of which show the reliability, generality, and scalability of the pro-

[☆]This work is partially supported by National Key R&D Program of China under the Grant No. 2017YFC0804002, National Science Foundation of China under the Grant No. 61761136008, Science and Technology Innovation Committee Foundation of Shenzhen under the Grant No. ZDSYS201703031748284, Shenzhen Peacock Plan under the Grant No. KQTD2016112514355531, Program for Guangdong Introducing Innovative and Entrepreneurial Teams under the Grant No. 2017ZT07X386.

*Corresponding author: Yuhui Shi.

Email addresses: 11860004@mail.sustech.edu.cn, 13278371@student.uts.edu.au (Lijun Sun), 11849557@mail.sustech.edu.cn (Chao Lyu), shiyh@sustech.edu.cn (Yuhui Shi)

posed CCPSO-R. Comparison with a representative dynamic path planning based algorithm Multi-Agent Real-Time Pursuit (MAPS) further shows the effectiveness of CCPSO-R. Finally, the codes of this paper are public available at: <https://github.com/LijunSun90/pursuitCCPSOR>.

Keywords: Swarm intelligence, Cooperative coevolution, Particle swarm optimization, Pursuit domain, Virtual robot

1. Introduction

The pursuit domain, or predator-prey problem is a classical and interesting research domain which acts as one of the widely used fundamental testbeds for coordination techniques since it was proposed by Benda et al. [1]. On one hand, its apparently simple problem setup and flexibility in approaches or concept evaluations lead to both its popularity and the toy domain impression. On the other hand, it is challenging and thus a good domain for the research of swarm intelligence emerged from the cooperation among robots or agents, which has drawn much attention of researchers on various versions of the pursuit domain.

At first, greedy coordination strategies were manually designed by Korf [2], part of which were improved by Haynes et al. [3]. After that, Haynes et al. [3–6] improved the pursuit performance using evolutionary algorithms, such as genetic programming (GP) [7], strongly typed genetic programming (STGP) [8], and cases learning methods successively. However, these methods cannot assure 100% capture. In 2000s, Undeger and Polat [9] treated the multi-agent dynamic pursuing problem in partially observable environments with obstacles as a dynamic path planning and task allocation problem and proposed the multi-agent real-time pursuit (MAPS) algorithm. Besides, much works have been done in the field of reinforcement learning (RL). For example, Ishiwaka et al. [10] investigated the mechanism of the emergence of the predators' cooperative behaviors aiming to capture the prey in the continuous world. Barrett et al. [11, 12] evaluated the designed single agent in the ad hoc teamwork and took the pursuit domain as one benchmark task. As researches going on, the capture

reliability and the efficiency of approaches have both been improved. A detailed survey on the pursuit domain can be found in [13].

In this paper, we deal with the dynamic pursuit domain problem with a scalable swarm of predator robots and types of the prey in bounded diagonal grid worlds. Different from prior work, this paper treats the pursuit domain as an optimization problem and proposes a particle swarm optimization (PSO) based cooperative coevolutionary (CC) algorithm, called CCPSO-R (R is for robots), where, to the best knowledge of authors, real and virtual robots coexist for the first time in an evolutionary algorithm (EA). In detail, we have n subpopulations, each of which has the same population size and evolves independently. The first individual of each subpopulation always corresponds to a unique real robot, which constitutes the swarm of cooperative real predator robots pursuing the prey. The rest are virtual robots, which are always deployed around their corresponding real robots, exploring the real robot’s vicinity in order to guide the real robot to a more advantageous position under the supervision of the fitness function defined on the pursuit task. Hence, in the view of the multi-agent system (MAS), these virtual robots can be seen as the action space for each real robot. Since the virtual robots only occupy part of a real predator’s vicinity, the exploration of virtual robots is actually a sampling rather than an exhausted exploration to a vicinity, which is guided by a proven efficiency swarm intelligence algorithm—PSO [14]. Therefore, the proposed CCPSO-R can be expected to be more efficient and effective.

In addition, the collision avoidance consideration among real robots is integrated into the fitness function design, which not only separates the robotic considerations from the EA itself and is thus different from the robotic PSO (RPSO) [15]—the PSO variant specially designed for robots, but also enhances the flexibility of the fitness function by modular design.

Furthermore, unlike previous incremental construction based EA methods and RL algorithms, the proposed CCPSO-R is actually an on-line algorithm which plans one step ahead for each robot and can reliably capture the prey even without the training and learning stage under the immediate guidance of the

fitness function. Meanwhile, similar to the common strategy in RL algorithms, the other real robots (agents in MAS) are treated as parts of the dynamic environment to the current robot without any central commander/controller.

The rest of the paper is organized as follows. First, the pursuit domain definition and details adopted here are explained in Section 2. Then the proposed CCPSO-R is described in Section 3. Experiments, comparisons, corresponding results, and discussions are presented in Section 4. Finally, conclusions and directions for future researches are given in Section 5.

2. The pursuit/predator-prey domain

Generally speaking, the pursuit domain problem can be considered as a game where predators try to capture the prey with or without coordination. However, as summarized in [13] and mentioned above, the pursuit domain has various versions depending on different combinations of its parameters, such as the type and size of the world, definition of the capture, team size, legal moves and move orders for the predators and prey, distance metric, etc. In many researches, a toroidal world is selected to simulate an infinite world, where a robot comes out of one edge will come in immediately from the opposite edge. However, this kind of world is not practical. As depicted in Figure 1a, if the




Figure 1: One trick of the coordination strategy in a toroidal world.

red pentagram is a linear prey which moves in a straight line towards north and

just escapes the nearly encirclement of the predators (blue squares), in the real infinite world, the predators will never catch the prey if they have the same speed. But in the toroidal world, if the predators move as shown in Figure 1b, they will capture the linear prey in the next step. Therefore, in this paper, rather than toroidal worlds, bounded grid worlds are selected, which can at least represent partially, although not all, the real world scenarios, such as an indoor room or an outdoor park with boundaries, etc.

Besides, as classified by Korf [2], the game with a discrete world (grid world here) that only allows horizontal and vertical, totally 4 directions movements, is called the orthogonal game, while the one which allows the horizontal, vertical and diagonal 8 directions move is called the diagonal game. Again, towards real applications, the diagonal game is more realistic [2] and thus one of the assumptions of this paper. In particular, no collisions are allowed and orthogonal obstacles will be considered when the real (predator or prey) robot moves diagonally, as illustrated in Figure 2. Under these assumptions, the capture can

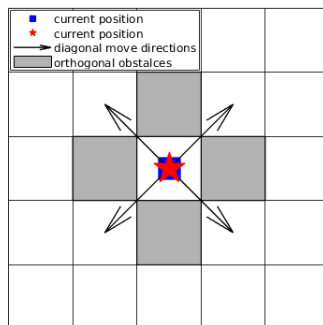


Figure 2: Illustration of the orthogonal obstacles.

be defined as that every available orthogonal neighbor of the prey robot has been occupied by a predator robot as shown in Figure 3. This may be different from the definitions of some research work, especially the RL algorithms or path planning based methods [9], where the capture is defined as that the position of the prey is occupied by a predator. Thus, the capture is more difficult here.

As for the other details of the pursuit game, the prey robot always moves




Figure 3: Illustration of the capture, taking 4 predator robots as an example.

first and then the predator robots move one by one in a fixed order, the natural priorities of which can make the collision avoidance control easier and more reliable. Besides, the position of any real (prey or predator) robot is visible to all the other robots. But there is no explicit communications, i.e. no explicit negotiations or coordination, among the robots. In other words, the coordination among real predator robots, or subpopulations, is implicit here. In addition, as will be seen later, no fixed behavior rules for the predator robots exist due to the fact that the evolution, or the one step ahead plan, of a predator robot in the dynamic environment is only guided by the fitness function.

3. Cooperative coevolution of real and virtual robots

In this paper, coevolved predators cooperate to encircle a prey, and the evaluation function is called the fitness function in EA [16], which is (functionally) identical to an objective function in the optimization field. So, the pursuit domain problem can be treated as an optimization problem in the sense that the goal is to improve the fitness of the pursuit process. Concretely, the optimization is conducted by a particle swarm optimization (PSO) based cooperative coevolutionary (CC) algorithm called CCPSO-R.

3.1. Fitness function

According to the capture definition in Section 2 and the task that a swarm of predator robots needs to encircle a prey robot, the fitness function should subject to the following metrics:

- *CLOSURE* $f_{closure}$: the prey robot should locate inside the convex hull of the predator robots positions;
- *SWARM EXPANSE* $f_{expanse}$: the swarm of predator robots should concentrate around the prey robot, i.e., a smaller swarm expanse of the predator robots is preferred;
- *UNIFORMITY* $f_{uniformity}$: the predator robots should distribute uniformly around the prey robot;
- *COLLISION AVOIDANCE* f_{repel} : collisions among real (predator/prey) robots are not allowed in the practical sense.

It is obvious that a single predator robot itself cannot form a solution. In CCPSO-R, a complete solution to the pursuit problem is composed by the positions of all the predator robots. However, before formulating the fitness function, a definition needs to be introduced first.

Definition of Convex Hull [17]: The *convex hull* of the point set P , denoted by $\text{conv}(P)$, is the intersection of all convex regions that contain P .

An intuitive illustration of this definition can be found in [17] as in Figure 4.

Further, we define the function:

$$inconv(p, \text{conv}(P)) \stackrel{\text{def}}{=} \begin{cases} 0, & \text{if point } p \text{ is in } \text{conv}(P) \\ 0.5, & \text{if point } p \text{ is on the edge of } \text{conv}(P) \\ 1, & \text{otherwise} \end{cases} \quad (1)$$

Hence, the fitness function for the j th ($j = 1, \dots, N_p$) individual (robot) in the i th ($i = 1, \dots, N_s$) subpopulation p_{robots}^{ij} is defined as

$$f^{ij} = f_{repel}^{ij} \cdot (f_{closure}^{ij} + f_{expanse}^{ij} + f_{uniformity}^{ij}) \quad (2)$$

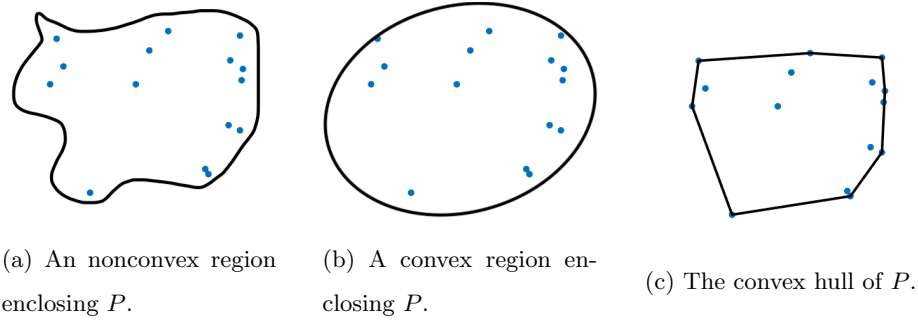


Figure 4: Illustration of the convex hull of the point set P [17].

where

$$f_{repel}^{ij} = \begin{cases} e^{-2 \cdot (NND^{ij} - D_{min})}, & \text{if } NND^{ij} < D_{min} \\ 1, & \text{else} \end{cases} \quad (3)$$

corresponds to the above *COLLISION AVOIDANCE* metric,

$$f_{closure}^{ij} = inconu(p_{prey}, conv(p_{robots}^{11}, \dots, p_{robots}^{ij}, \dots, p_{robots}^{N_s 1})) \quad (4)$$

corresponds to the above *CLOSURE* metric,

$$f_{expanse}^{ij} = \frac{1}{N_s} \left(\sum_{k=1, k \neq i}^{N_s} |p_{robots}^{k1} - p_{prey}| + |p_{robots}^{ij} - p_{prey}| \right) \quad (5)$$

corresponds to the above *SWARM EXPANSE* metric, and

$$f_{uniformity}^{ij} = std \left(\begin{bmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{bmatrix} \right) \quad (6)$$

corresponds to the above *UNIFORMITY* metric.

In the formulas, NND^{ij} is the nearest neighbor distance, i.e., the minimum of the pairwise Euclidean distances between the j th individual in the i th subpopulation and all the real predator robots in the other subpopulations; D_{min} is a specified secure distance for collision avoidance; p_{prey} is the position of the prey robot; p_{robots}^{ij} is the position of the j th robot in the i th subpopulation; $std(\cdot)$ stands for the standard deviation function; and N_{kh} ($k = 1, 2; h = 1, 2$) is the counts of the real predator robots in the (k, h) -th bin out of the overall

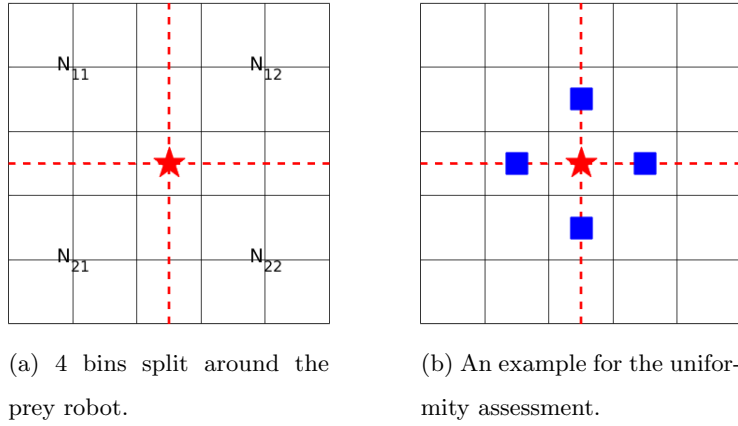


Figure 5: Illustration of the uniformity assessment.

4 bins split by the horizontal and vertical lines which intersect at the position of the prey robot, as shown in Figure 5. Note that, the number of the real predator robots on the split lines is divided by 2 and equally assigned to the two adjacent bins. Hence, $N_{11} = N_{12} = N_{21} = N_{22} = 1$ for the example in Figure 5b. However, the formula (6) cannot always give the objective uniformity

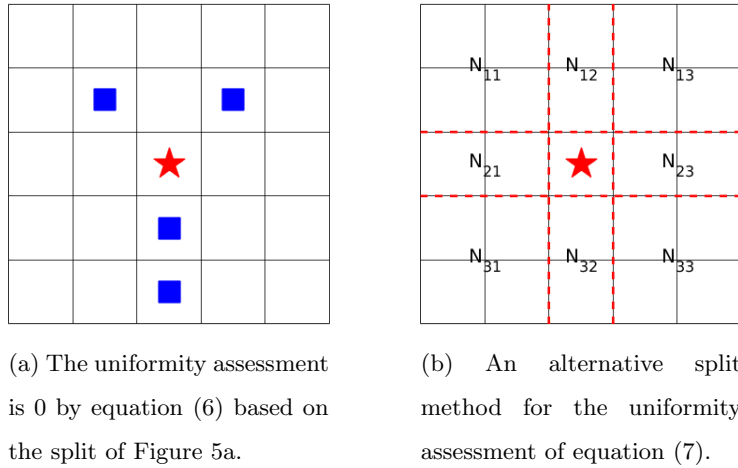


Figure 6: Illustration of the alternative uniformity assessment.

assessment that is consistent with a human's subjective judgment, as the deadlock phenomenon shown in Figure 6a. In this scenario, the prey robot always

keeps still in the center of the map, while the predator robots start to encircle the prey from randomly generated initial positions and stop forever since the game state shown in Figure 6a, which is obviously not the expected capture state. So, the deadlock phenomenon is a game state where the pursuit task hasn't been accomplished but all the predator robots stop forever as if they are locked. One reason of the deadlock phenomenon is the fitness function wrongly evaluates an intermediate game state as fittest, i.e., the task is finished. Therefore, to design a better fitness function, in the uniformity assessment formula, an alternative space split strategy is performed as shown in Figure 6b, and the following uniformity assessment will replace equation (6) in such situations:

$$f_{uniformity}^{ij} = std([N_{12}, N_{21}, N_{23}, N_{32}]) + std([N_{11}, N_{13}, N_{31}, N_{33}]), \quad (7)$$

the first and second part of which are the axial and diagonal uniformity assessments, respectively.

To be clearer, the fitness evaluation of p_{robots}^{ij} , i.e., the j th ($j = 1, \dots, N_p$) individual (robot) in the i th ($i = 1, \dots, N_s$) subpopulation, is illustrated in Figure 7.




Figure 7: Illustration of the fitness evaluation for p_{robots}^{ij} .

3.2. The proposed CCPSO-R algorithm

In CCPSO-R, there are N_s independently evolved subpopulations with subpopulation size N_p , and the first individual of each subpopulation represents a

unique real robot while the others represent virtual robots. All the real robots consist of the predator robots swarm which actually pursue the prey robot in the grid world, while the virtual robots are to explore the vicinity of the corresponding real predator robot in its subpopulation and guide the predator robot to a better position. So in this sense, virtual robots can be seen as the action space of the corresponding real predator robot. The real predator robot chooses its locally optimal action, but in terms of the global benefit of the whole swarm of predators. That is, the evaluation of a robot position is conducted by considering the rest real predator robots positions in the other subpopulations. Since the proposed algorithm works in the modes of cooperative coevolutionary algorithms (CCEAs), it is called the cooperative coevolutionary PSO for robots (CCPSO-R), as illustrated in Algorithm 1, which will be explained in detail from 3 aspects: the update rules, the fitness evaluation, and the diversity maintenance mechanism in the following.

3.2.1. Update rules

Two update rules are designed separately for virtual and real robots:

1. For a virtual robot j ($j \in \{2, \dots, N_p\}$), the PSO update rules are as follows:

$$v_{robots}^{ij} = nnd(w \cdot v_{robots}^{ij} + c_1 \cdot r_1 \cdot (p_{robots}^{ij} - p_{robots}^{ij}) + c_2 \cdot r_2 \cdot (pg_{robots}^i - p_{robots}^{ij})) \quad (8)$$

$$p_{robots}^{ij} = nbn((p_{robots}^{ij} + v_{robots}^{ij}), p_{robots}^{i1}) \quad (9)$$

where

$$nnd(v) = \arg \min_{p_n \in SN} |\angle p_n - \angle v| \quad (10)$$

and

$$SN = \{(1, 0), (1, 1), (0, 1), (-1, 1), (-1, 0), (-1, -1), (0, -1), (1, -1)\}. \quad (11)$$

$nnd(v)$ outputs one of the 8 unit vectors in SN which has the minimum angle distance with the input velocity v . By using the function $nnd(\cdot)$, every robot

Algorithm 1: CCPSO-R

```
1 Initialization
2 while the prey is not captured and time limit is not reached do
3   for each subpopulation do
4     Re-evaluate the subpopulation due to environmental changes
5     for each virtual robot do
6       Update its velocity and position using (8) and (9)
7       Evaluate the fitness together with the rest real predator
8         robots
9       if unique virtual robots  $< T_v$  then
10        Re-initiate and re-evaluate the virtual robots
11      Update the velocity and position of the real predator robot using
12        (13) and (14)
13      Evaluate the fitness of the real predator robot together with the
14        rest real predator robots
15      if the real predator robot becomes the global best then
16        Re-initiate and re-evaluate the virtual robots
17      else if the predator robot gets trapped in a deadlock then
18        Add a random noise to the real predator robot position
19        Re-evaluate the whole population
20    if the real predator robots swarm get trapped in a local optimum
21      then
22        Add random noises to all the real predator robots positions
23        Re-evaluate the whole population
```

can only move one step by one step. In this way, unlike the multi-steps case in a general PSO, the path planning and the worry about collisions in the half way to a destination are not ever necessary. v_{robots}^{ij} is the velocity for the j th individual (robot) in the i th subpopulation which has the position p_{robots}^{ij} . In addition,

p_{robots}^{ij} is the individual historical best position for the j th individual (robot) in the i th subpopulation, while pg_{robots}^i is the global best position of the i th subpopulation. The coefficient $w \in R$ is called the inertia weight, $c_1, c_2 \in R^+$, and r_1, r_2 are uniformly distributed random numbers in the range of $(0, 1)$. Besides,

$$nbn(p_{robots}^{ij}, p_{robots}^{i1}) = \begin{cases} \arg \min_{p_b^{i1}} |\angle(p_b^{i1} - p_{robots}^{i1}) - \angle(p_{robots}^{ij} - p_{robots}^{i1})|, & \text{if } p_{robots}^{ij} \text{ is out of the vicinity of } p_{robots}^{i1} \\ p_{robots}^{ij}, & \text{otherwise} \end{cases} \quad (12)$$

is designed to output the nearest boundary neighbor p_b^{i1} in the constrained vicinity of the real predator robot p_{robots}^{i1} . This function is illustrated in Figure 8, where the constrained vicinity of p_{robots}^{i1} is shown in a dashed square, which is determined as the minimum one that can accommodate the specified number of virtual robots. Note that, the $nbn(p_{robots}^{ij}, p_{robots}^{i1})$ function in Equation (9) is very important because it can assure all the virtual robots p_{robots}^{ij} ($j \geq 2$) are in the constrained vicinity of the real predator robot p_{robots}^{i1} , without which the subpopulation may lose the vicinity exploring capability for the real predator robot.




Figure 8: Illustration of the function $nbn(p_{robots}^{ij}, p_{robots}^{i1})$ in Equation (12).

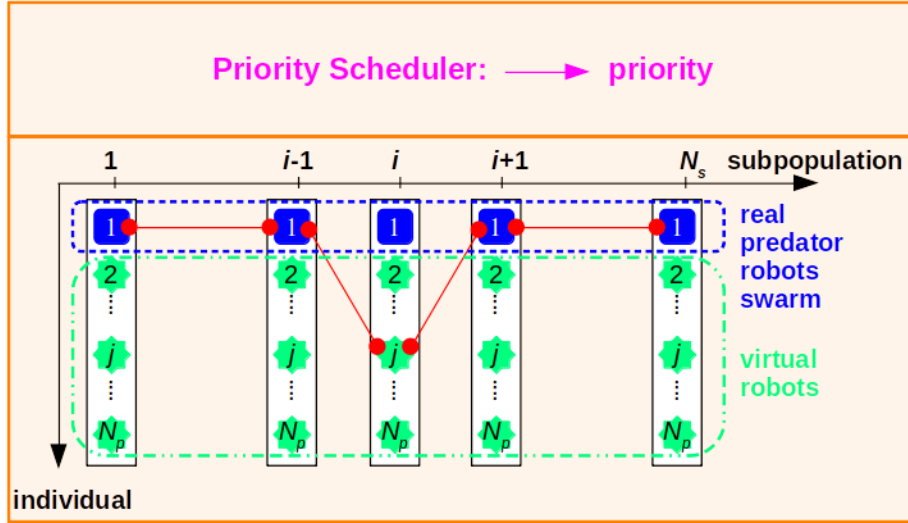


Figure 9: Illustration of the priority scheduler for the CC based fitness evaluation.

2. For the real robot ($j = 1$), the PSO update rules are as follows:

$$v_{robots}^{i1} = nnd(pg_{robots}^i - p_{robots}^{i1}) \quad (13)$$

$$p_{robots}^{i1} = p_{robots}^{i1} + v_{robots}^{i1} \quad (14)$$

So, the real predator robot does not need to perform any exploring task, but just quickly becomes the global best in its subpopulation.

To summarize, by utilizing different optimization mechanisms for different kinds of robots, virtual robots are responsible for exploring and finding potential better positions in the vicinity of the real predator robot, while the real predator robot in each subpopulation just makes use of the achievements of the virtual robots and becomes the global best.

3.2.2. Fitness evaluation

From the practical point of view, no collisions of any two real robots are allowed. Since we have totally N_s subpopulations in the cooperative coevolution population, a priority scheduler is used to coordinate among them, as shown in Figure 9. In particular, a priority scheduler will decide the order of movements among real robots, which can be the evolving order of subpopulations

in asynchronous situations, and can also be used to coordinate, for example, two real robots when they both want to go to the same position in synchronous situations.

To be as simple as possible, here we let the priorities be in consistent with the indexes of the subpopulations. In other words, after the prey robot moves, the subpopulations evolve one-by-one and the newly updated real predator robot is counted into the dynamics of the environment for the fitness evaluations to the subsequent subpopulations. So, if $k > h$, the predator p_{robots}^{k1} always moves ahead of the predator p_{robots}^{h1} .

To evaluate the fitness of the j th individual in the i th subpopulation, a complete solution should be first composed by replacing the i th real predator robot with p_{robots}^{ij} from the real predator robots swarm:

$$\left[p_{robots}^{11}, \dots, p_{robots}^{(i-1)1}, p_{robots}^{ij}, p_{robots}^{(i+1)1}, \dots, p_{robots}^{N_s 1} \right].$$

Then, the fitness of a robot can be evaluated by equation (2), as shown in Figure 7.

3.2.3. Diversity maintenance mechanism

When a swarm intelligence algorithm converges, all individuals may be attracted to the same position, no matter it is the global or local optimum. However, for the pursuit case here, the convergence of virtual robots in a subpopulation brings the disadvantage that the capability of exploring potentially better positions is getting worse. Therefore, if the number of unique virtual robots in a subpopulation is defined as the subpopulation diversity, the diversity of each subpopulation must be maintained to keep its exploring capability. Besides, due to the existence of unexpected deadlocks, suitable strategies should be integrated in the coordination algorithm to deal with such problems.

Based on the above ideas, we propose the diversity maintenance mechanisms which are performed as follows:

- Update the population in each generation based on the scheme that the fitness of the newly generated individual is not worse than its parental

robot, which will guide the robot to explore more positions without harm to the fitness.

- Redistribute the virtual robots once the number of unique virtual robots positions in a subpopulation decreases below a threshold T_v , i.e., the subpopulation converges. That is, the subpopulation has found better solutions and all robots are attracted to the global best. In this situation, virtual robots should be redistributed to the space for better exploration. This strategy corresponds to the line 8-9 in Algorithm 1.
- Redistribute virtual robots once the real predator robot becomes the global best in the subpopulation. Because the role of virtual robots is to help the corresponding real predator robot to find better positions, once this real predator robot becomes the global best in its subpopulation, the object of virtual robots is reached and they should be redistributed to the space to find potential better positions for the real predator robot. This strategy corresponds to the line 12-13 in Algorithm 1.
- Add a random noise to the position of the real predator robot if it is not the global best in its subpopulation but abnormally keeps stills for a long time, in which it must have gotten stuck in a deadlock. This strategy corresponds to the line 14-16 in Algorithm 1.
- Add random noise to the positions of all the real predator robots if they converge when the prey robot has not been captured, the situation of which can be seen as that the swarm of predator robots gets trapped in a local optimum. This strategy corresponds to the line 17-19 in Algorithm 1.

4. Experiments

In this section, two different experiments are presented. Experiment 1 is conducted in a 30×30 grid world to verify the performance of the proposed CCPSO-R. Experiment 2 is to compare CCPSO-R with a representative dynamic path

planning based pursuit algorithm MAPS [9]. From the experimental results , pros and cons of two different strategies can be seen in spite that CCPSO-R and MAPS are originally designed towards different capture definitions.

In particular, to verify the generality of algorithms, four types of preys are implemented. The prey robot initially locates in the center of the world, but behaves differently according to its type defined as follows:

- STILL PREY: the still prey keeps still in its initial position forever.
- RANDOM PREY: the random prey randomly moves to a next position according to the uniform distribution.
- LINEAR PREY: the linear prey initially chooses one of the 8 directions in which the number of predator robots is minimum, and moves in that direction in a straight line since then. Only when the prey locates on edges of the map, it will re-calculate a new direction according to the same criterion. However, when the way of the linear prey is blocked by a predator, it cannot move any more but only wait for the other predator robots coming to encircle it.
- SMARTER LINEAR PREY: the smarter linear prey, represented as `linear_smart`, is very similar to the linear prey. The only difference is that when its way is blocked by a predator robot, it moves to an unoccupied neighbor which has the minimum angle distance with its current direction and then it continues its movement in its previous direction if there are no obstacles.

From the above descriptions, the capabilities of the preys and the difficulties of encircling preys can be intuitively ranked as “still prey < random prey < linear prey < linear_smart prey”, which will be further verified by the following experiments.

4.1. Experiment 1

4.1.1. Experimental setup

To verify the scalability of CCPSO-R, various sizes of the swarm of predators, i.e., 4, 8, 12, 16 and 24, are used, from which we can expect the advantages originated from the swarm intelligence of the swarm of predator robots.

The other implementation details are as follows: the initial real predator robots are deployed randomly in the whole grid world without overlapping; the population size of each subpopulation is 20; the prey robot moves in 90% of the time ensuring that predators move faster or a longer distance than the prey; in equation (3) $D_{min} = 1$ which is the minimum secure distance between two robots; in equation (9) the parameters $w = 1, c_1 = c_2 = 2$ which are set as recommended in [14]; T_v is 9 in the line 8 of Algorithm 1 which is the number of grids for a 3×3 vicinity; when the real robot is not the global best in its subpopulation but keeps still over 5 iterations we say that it gets trapped in a deadlock which corresponds to the line 14 of Algorithm 1; when the swarm of predator robots keeps still over 10 iterations we say the swarm converged, and if the swarm has converged but the prey hasn't been captured we say that the swarm gets trapped in a local optimum which corresponds to the line 17 of Algorithm 1.

In addition, for environmental changes such as real predator robot position change in other subpopulations, the current subpopulation needs to be re-evaluated as shown in the line 4 of Algorithm 1, where the individual historical best position pi_{robots}^{ij} will not be inherited, and the global best pg_{robots}^i will be re-calculated. This is because, although the experimental results of inheriting and not inheriting the individual historical memory pi_{robots}^{ij} differ, it is hard to select either one due to their competitive performances.

As for the performance metrics, we use the number of successful captures, the average number of moves to capture the prey, and their standard deviations over 100 randomly generated test cases given the maximum 1000 time steps, the random seeds of which are set from 1 to 100.

4.1.2. Simulation results

Table 1: Number of captures, average number of moves, and their standard deviations to capture different preys with various number of predators out of 100 test cases.

No. of predators	Metrics	Prey			
		Still	Random	Linear	Linear_Smart
4	No. of captures	100	100	100	100
	Avg. of moves	30.450	49.840	46.900	204.060
	Std. of moves	19.943	36.867	31.886	198.927
8	No. of captures	100	100	100	100
	Avg. of moves	22.220	33.780	42.240	121.820
	Std. of moves	13.384	23.039	46.583	111.922
12	No. of captures	100	100	100	100
	Avg. of moves	20.470	24.520	30.190	76.780
	Std. of moves	11.364	13.414	24.943	72.839
16	No. of captures	100	100	100	100
	Avg. of moves	17.360	18.360	25.620	49.850
	Std. of moves	9.648	11.277	23.560	50.048
24	No. of captures	100	100	100	100
	Avg. of moves	15.060	14.060	19.670	35.400
	Std. of moves	10.688	6.151	21.879	32.588

The simulation results are summarized in Table 1, from which it can be seen that CCPSO-R is reliable with the capture rate being 100% in a limited time, no matter what type of the prey it is. As expected before, to a swarm of predator robots, the difficulties, in terms of the average number of moves, to capture each type of prey can be generally ranked as “still prey < random prey < linear prey < linear.smart prey”, which can be seen more clearly from Figure 10. This conclusion is in consistent with the common opinion in literature (such as [5] and [13]) that compared with the random prey, the straight line moving prey

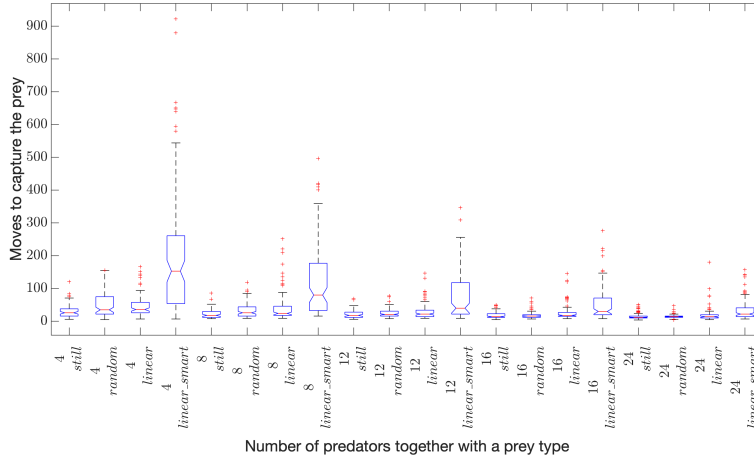


Figure 10: Box plot of the moves to capture a specific prey with a specific number of predator robots.

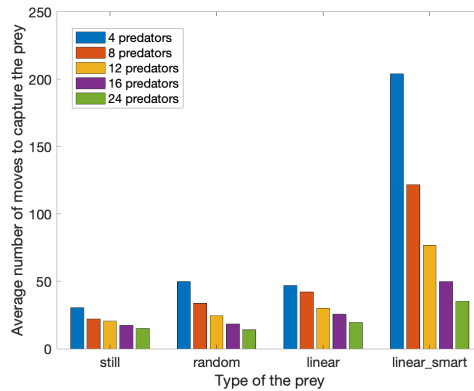


Figure 11: Bar graph of the average moves to capture a specific prey with a specific number of predator robots.

is more effective because it breaks the movement locality. Hence, the straight line moving prey is more difficult to be captured, which leads to the low capture rates in previous work, such as the manually designed methods [2, 5], EA based method [5] and the case learning method [6].

In addition, we show the data of Table 1 in the manner of Figure 11, from which an evident fact can be found that the more predator robots the more

efficient the pursuit is. Besides, from the decreasing standard deviations as more real predator robots are involved, as shown in Figure 10, it can be concluded that with the swarm size of the predator robots gets larger, the pursuit performance is getting more and more stable and robust.

To give a more intuitive impression of the pursuit process, several representative episodes taken from an experiment against the linear_smart prey are displayed in Figure 12.

4.1.3. Discussions

The results achieved in Section 4.1.2 can be explained from the algorithm's point of view. First, as we treat the pursuit domain as an optimization problem, as long as the designed fitness function (or objective function) properly models the investigated problem, minimizing the fitness function will lead predators to a successful capture. Second, on one hand, every step of a real predator robot is greedy since it moves to the best virtual robot position in its subpopulation; on the other hand, each step of a real predator is not totally greedy since the virtual robots exploration in its vicinity is not exhausted. So, predators may eventually capture the prey but the process may be slow.

Besides, in some of the past work, such as RL and path planning approaches, the capture is defined differently as that the prey position is occupied by a predator, the further idealization and simplification of which result in an unpractical problem setup for robots applications. This paper adopts another conventional definition that the prey is encircled by predators such that it cannot move any more, which considers both the collision avoidance and the safety of robots from the practical point of view. However, to further validate the effectiveness of the proposed CCPSO-R, we will do the comparison in the next Sub-section.

4.2. Experiment 2

For pursuit domain problems, path planning and task allocation based strategies are intuitive and may be the first solution that comes to one's mind. Therefore, in this section, we compare the proposed CCPSO-R with a representative

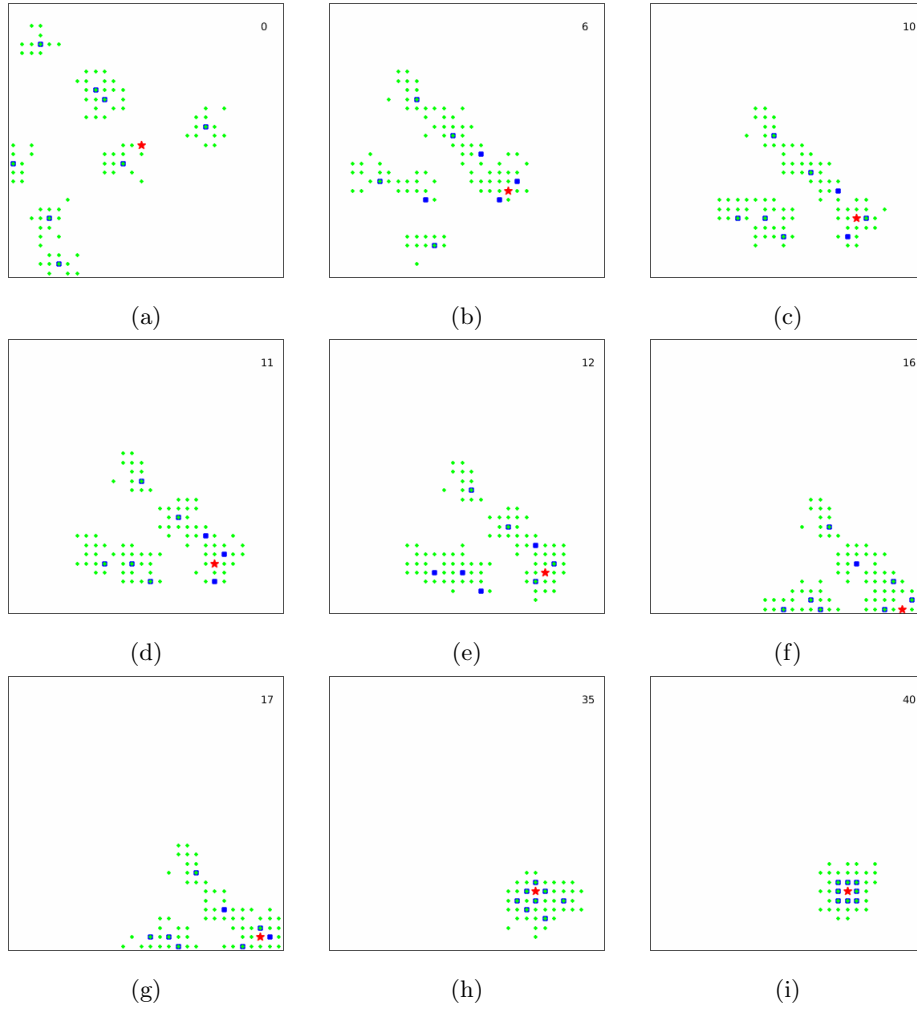


Figure 12: Illustration of the pursuit process, taking the pursuit of a linear_smart prey as an example. (a) is the initialization. From (a)-(b), the prey moves in the southeast direction in a straight line. In (c), the prey encounters an orthogonal real predator robot. So, in (d), the prey moves to a nearest unoccupied neighbor in the south. After that, from (d) to (e), the prey continues to move in its previous straight line direction. Until (f), the prey reaches an edge. So, in (g), the prey re-selects the north as its new escape direction. In (h), the prey is captured. And in (i), the predator robots swarm converge.

dynamic path planning based algorithm named MAPS [9].

4.2.1. Experimental setup

All the experimental setups are the same as the Experiment 1 except the number of predator robots, the grid world sizes, and the way we calculate the capture status. For fair comparisons, we modify the capture definition to the one adopted by MAPS that the position of the prey is occupied by any predator. In particular, we know that in CCPSO-R, when predators cooperate to encircle the prey, there must be at least one such moment that the prey is adjacent to a predator and they get common neighbors as illustrated in Figure 13. Because diagonal obstacles are considered in our collision avoidance design, the number of the common neighbors is at most 2 except the prey's own position. If the prey keeps still or moves to anyone of the common neighbors the next moment, no matter it is compelled by other coordinated predators or just due to the simplicity of itself, the adjacent predator can definitely occupy the prey's new location since predators always move after the prey. We can then say that the prey is captured at this moment.

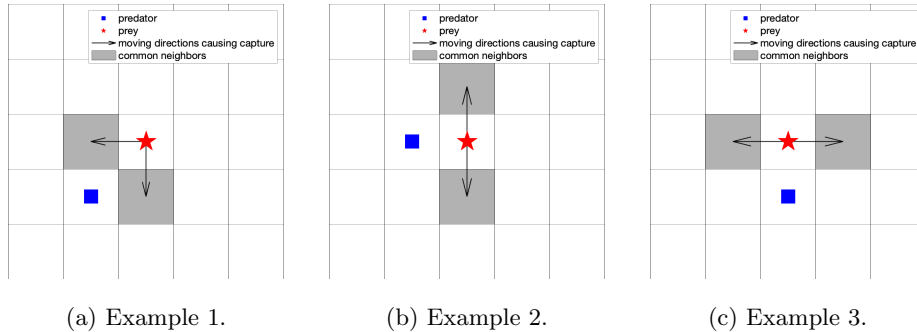


Figure 13: Illustration of the modified capture status of CCPSO-R used in the comparison with MAPS.

4.2.2. Simulation results

We run the two algorithms MAPS and CCPSO-R in the same randomly generated scenarios where the initial predators and prey positions are the same.

Table 2: Comparison results: Average number of moves and their standard deviations to capture different preys with various number of predators in different sizes of grid worlds out of 100 test cases.

Metrics		Avg. of moves		Std. of moves		
Number of predators		4	8	4	8	
Size of the grid map: 30×30						
Prey	Still	MAPS	7.09	4.87	3.777	2.751
		CCPSO-R	6.5	4.66	2.819	2.271
	Random	MAPS	9.01	6.78	6.973	5.868
		CCPSO-R	8.3	6.04	3.335	2.828
	Linear	MAPS	17.89	10.55	6.977	8.184
		CCPSO-R	16.9	7.71	8.487	5.695
	Linear	MAPS	17.84	10.55	7.102	8.211
	Smart	CCPSO-R	16.92	7.75	8.474	5.723
Size of the grid map: 150×150						
Prey	Still	MAPS	36.6	26.16	18.229	14.633
		CCPSO-R	33.37	23.07	13.9	10.992
	Random	MAPS	38.62	28.55	18.968	16.516
		CCPSO-R	35.84	26.07	13.654	11.401
	Linear	MAPS	99.33	64.29	44.576	37.918
		CCPSO-R	82.75	41.32	30.274	29.684
	Linear	MAPS	99.19	64.37	44.849	37.814
	Smart	CCPSO-R	82.84	41.44	30.192	29.798

To see whether the map size influences the comparison results much, a bigger grid world size 150×150 together with a smaller grid map size 30×30 are applied with two predator swarm sizes 4 and 8. Experimental results are summarized in Table 2. Since both algorithms capture the prey 100%, we do not list the “No. of captures” metric. It is not hard to understand that generally with the increase of the predators number, less moves with smaller standard deviations

Table 3: Statistical test for CCPSO-R vs. MAPS using the t-test at the 5% significance level.

World size		30 × 30		150 × 150	
No. of predators	Prey	p-value	Significant	p-value	Significant
4	Still	0.013848	Yes	0.0062804	Yes
	Random	0.13224	No	0.02533	Yes
	Linear	0.10273	No	7.3641e-05	Yes
	Linear Smart	0.11762	No	9.8947e-05	Yes
8	Still	0.13084	No	0.0017754	Yes
	Random	0.11016	No	0.027238	Yes
	Linear	0.00018443	Yes	1.3396e-11	Yes
	Linear Smart	0.00020334	Yes	1.1589e-11	Yes

are needed for a capture; and with the increase of the grid map size, values in these two metrics increase accordingly.

To further validate the significance of the comparison results in Table 2, t-tests are conducted at the 5% significance level, as shown in Table 3. It can be seen that CCPSO-R significantly outperforms MAPS in 11 out of the total 16 experimental scenarios where more predators are involved with smarter preys in bigger worlds.

4.2.3. Discussions

The comparison results show that although CCPSO-R is designed for the coordinated encirclement of the prey until every available orthogonal neighbor of the prey has been occupied by a predator, CCPSO-R also performs well in the position occupying based capture definition, which proves the effectiveness of the cooperative coevolutionary based coordination strategy.

In addition, compared with just occupying the same position of the prey with

only one predator, occupying all the orthogonal neighbors of the prey simultaneously as uniformly as possible is more complicated. Therefore, the proposed CCPSO-R can accomplish more complicated coordination tasks compared to MAPS.

On the other hand, as presented in [9] that MAPS is a real-time pursuit algorithm, its MATLAB version, which is rewritten by us from its original C++ codes, still runs faster than CCPSO-R. Because the calculation of the average number of moves per second over all the test cases has a high requirement on the runtime environment for fair comparison, we do not list this metric here. But we can still get the following conclusions. MAPS is faster. However, one problem of it is that its performance is constrained by the number of predators. Because one important step in MAPS is to assign the possible escape directions of the prey to every predator optimally by iterating every possible assignments, the combinatorial number of which is $(n - 1)!$ where n is the number of predators. With the increase of n , this combinatorial number will increase very fast and it is becoming less practical to get all the permutations at once, which also brings more burden to the memory. This is also the reason that we only compare the simulation results with up to 8 predators. But, for CCPSO-R, although it is not as efficient as MAPS, its scalability on the predator swarm size is much better.

Finally, despite the fact that both MAPS and CCPSO-R adopt the same sequential movement strategy that the prey always moves first and then the predators move one by one, another loop is embedded in the current CCPSO-R implementation that each real predator robot position can only be updated when all its corresponding virtual robots have been updated sequentially, as shown in line 5-10 of Algorithm 1. It is known that embedded loops are generally not expected for an efficient algorithm. So, a parallel update and evaluation for the virtual robots should alleviate the efficiency constrain of CCPSO-R which will be investigated in future work.

5. Conclusions

This paper treated the pursuit domain as an optimization problem and presented the cooperative coevolutionary algorithm—CCPSO-R, which, for the first time, introduces the combination of the real robots and virtual robots into the correspondences between the individual representation of an EA and the robots in an application. Before the work in this paper, an individual in an EA will be assigned to a real robot. However, in the proposed CCPSO-R algorithm, only the first individual in each subpopulation corresponds to a real robot, while the rest individuals are all virtual robots, who act as a kind of action space for real robots by sampling and exploring their vicinities.

Besides, it should be noted that there are no fixed behavior rules for the swarm of predator robots. Instead, the swarm of robots is guided directly by the fitness function, which is designed in a modular manner by incorporating very limited domain knowledge. As one module, the collision avoidance consideration is integrated in the fitness function, which itself is another fitness function for repelling and can be versatile by tuning its parameter D_{min} . If the $D_{min} = 1$, as it is in this paper, the robot swarm can capture the prey while moving without collisions.

Finally, we tested the performance—the generality, stability and scalability of the proposed CCPSO-R with four types of preys—the still prey, the random prey, the linear prey, and the linear smart prey. Experimental results have been summarized based on 100 randomly generated test cases whose random seeds are set as 1-100 for their reproducibility. Based on these experiments, it can be concluded that the proposed CCPSO-R can always capture the prey stably and no additional modifications are needed under different scenarios. In addition, a comparison with a representative dynamic path planning and task allocation based algorithm MAPS has also been conducted. Experimental results further prove the outstanding performance of the proposed CCPSO-R.

However, to be simple, the coordination priority scheduler was designed based on the subpopulation indexes, which indicates that the real predator

robots move in a fixed sequential order. This may be unreasonable when it is better to firstly move one specific predator which blocks others' ways. In addition, predators move sequentially, rather than synchronously, will deteriorate the pursuit efficiency when the swarm of predators gets larger. Therefore, three works need to be done in future: one is to study the memory inheritance strategy in dynamic optimization problems as mentioned in Section 4.1.1; one is to implement the parallel update and evaluation for the virtual robots; one is to improve the coordination scheduler towards the synchronous cooperation based on parallel computing by learning from experiences.

Acknowledgements

We would like to thank Dr. Cagatay Undeger for sharing their codes with us which made our comparison work feasible. Additionally, Professor Chin-Teng Lin helped us a lot during this research work, and Qiqi Duan gave kindly advice. Finally, we want to give our deeply appreciation to all the anonymous reviewers for their constructive comments and suggestions.

References

- [1] M. Benda, V. Jagannathan, R. Dodhiawala, On optimal cooperation of knowledge sources-an empirical investigation, Tech. rep., BCS-G2010-28, Boeing Advanced Technology Center, Boeing Computing Services, Seattle, Washington (1986).
- [2] R. E. Korf, A simple solution to pursuit games, in: Working Papers of The 11th International Workshop on Distributed Artificial Intelligence, 1992, pp. 183–194.
- [3] T. Haynes, R. L. Wainwright, S. Sen, Evolving cooperation strategies., in: ICMAS, 1995, p. 450.

- [4] T. Haynes, R. L. Wainwright, S. Sen, D. A. Schoenfeld, Strongly typed genetic programming in evolving cooperation strategies., in: ICGA, Vol. 95, 1995, pp. 271–278.
- [5] T. Haynes, S. Sen, Evolving behavioral strategies in predators and prey, in: G. Weiß, S. Sen (Eds.), *Adaption and Learning in Multi-Agent Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, 1996, pp. 113–126.
- [6] T. HAYNES, S. SEN, Learning cases to resolve conflicts and improve group behavior, *International Journal of Human-Computer Studies* 48 (1) (1998) 31 – 49. doi:<https://doi.org/10.1006/ijhc.1997.0159>.
- [7] J. R. Koza, *Genetic programming*, MIT Press, Cambridge, MA, 1992.
- [8] D. J. Montana, Strongly typed genetic programming, *Evolutionary computation* 3 (2) (1995) 199–230.
- [9] C. Undeger, F. Polat, Multi-agent real-time pursuit, *Autonomous Agents and Multi-Agent Systems* 21 (1) (2010) 69–107. doi:[10.1007/s10458-009-9102-0](https://doi.org/10.1007/s10458-009-9102-0).
- [10] Y. Ishiwaka, T. Sato, Y. Kakazu, An approach to the pursuit problem on a heterogeneous multiagent system using reinforcement learning, *Robotics and Autonomous Systems* 43 (4) (2003) 245 – 256. doi:[https://doi.org/10.1016/S0921-8890\(03\)00040-X](https://doi.org/10.1016/S0921-8890(03)00040-X).
- [11] S. Barrett, P. Stone, S. Kraus, Empirical evaluation of ad hoc teamwork in the pursuit domain, in: *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS '11*, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2011, pp. 567–574.
- [12] S. Barrett, A. Rosenfeld, S. Kraus, P. Stone, Making friends on the fly: Cooperating with new teammates, *Artificial Intelligence* 242 (2017) 132 – 171. doi:<https://doi.org/10.1016/j.artint.2016.10.005>.

- [13] P. Stone, M. Veloso, Multiagent systems: A survey from a machine learning perspective, *Autonomous Robots* 8 (3) (2000) 345–383. doi:10.1023/A:1008942012299.
- [14] Y. Shi, R. Eberhart, A modified particle swarm optimizer, in: 1998 IEEE International Conference on Evolutionary Computation Proceedings. IEEE World Congress on Computational Intelligence (Cat. No.98TH8360), 1998, pp. 69–73. doi:10.1109/ICEC.1998.699146.
- [15] M. S. Couceiro, R. P. Rocha, N. M. F. Ferreira, A novel multi-robot exploration approach based on particle swarm optimization algorithms, in: 2011 IEEE International Symposium on Safety, Security, and Rescue Robotics, 2011, pp. 327–332. doi:10.1109/SSRR.2011.6106751.
- [16] A. Eiben, J. Smith, Introduction to evolutionary computing, Springer, 2015.
- [17] S. L. Devadoss, J. O’Rourke, Discrete and computational geometry, Princeton University Press, 2011.